

CLAIMS

- 1 1. In a multi-slice network processor system comprising a plurality of processing slice
2 modules, each module processing and storing a slice of packet data, a method for processing a
3 packet in packet slices for transfer over a network interface comprising:
4 prepending a system header to the packet, the system header providing information for
5 use by the multi-slice system;
6 assigning a packet identifier to the packet;
7 segmenting data of the packet into cells, the data including both header and body data for
8 the packet;
9 generating cell descriptive information for each cell, the cell descriptive information
10 including the packet identifier, and a packet position indicator indicating an order position of
11 data of the cell with respect to the packet; and
12 delivering one or more cells of the packet to one or more processing slice modules based
13 upon load balancing criteria.
- 1 2. The method of claim 1 wherein load balancing criteria includes that no load balancing is
2 in effect.
- 1 3. The method of claim 1 wherein the packet identifier is a sequence number representing
2 an order of the packet in a communications flow and further comprising assigning a
3 communications flow indicator to the cell descriptive information of each cell of the packet.
- 1 4. The method of claim 1 wherein the cell descriptive information further comprises a slice
2 position indicator indicating an order position of the data of the cell with respect to a slice of data
3 of the packet.

1 5. The method of claim 3 further comprising delivering body data of the packet to one or
2 more of the processing slices ahead of the header data of the packet.

1 6. The method of claim 4 further comprising:
2 performing lookup functions for each slice of data;
3 determining a size of data change in header data; and
4 communicating the size of data change to a queue manager via an indicator in the system
5 header.

1 7. The method of claim 4 further comprising :
2 storing one or more cells in a buffer in the packet slice; and
3 generating a buffer correlation data structure correlating the buffers of the packet slice.

1 8. The method of claim 7 further comprising:
2 generating a slice correlation data structure for the packet based upon a packet reference
3 pointing to the buffer of the packet slice including the first cell of the packet, and the packet
4 identifier in each cell's descriptive information.

1 9. The method of claim 7 further comprising:
2 maintaining an independent set of upper bits of a sequence number for each
3 communication flow; and
4 responsive to detecting one of the processing slices delivering a sequence number that is
5 smaller in value than an immediately preceding sequence value for the same slice, incrementing
6 the independent set of upper bits for the respective communication flow, concatenating the set of
7 upper bits with a set of bits of the sequence number into an index, indexing into a re-sequencing
8 buffer space of sufficient depth to cover a slice-to-slice skew case based on the index, and
9 resequencing the packet into its sequence order position.

- 1 10. The method of claim 7 further comprising:
2 generating a slice correlation data structure for the packet including a packet reference
3 pointing to the buffer of the packet slice including the first cell of the packet, and a respective
4 buffer indicator for the buffer in each packet slice storing the first cell in the slice for the packet;
5 and
6 entering the slice correlation data structure as a single queue entry into a queue.
- 1 11. The method of claim 7 wherein the network interface is a switch fabric and further
2 comprising determining a destination slice across the switch fabric for each packet slice in
3 accordance with load balancing criteria.
- 1 12. The method of claim 11 further comprising:
2 for a received packet from the switch fabric, storing each cell of each packet slice of the
3 received packet, each cell including descriptive information, in the processing slice identified in
4 a destination slice indicator of the descriptive information.
- 1 13. The method of claim 12 further comprising sending an enqueue message for each packet
2 slice identifying a storage location of the first cell of the slice.
- 1 14. The method of claim 13 further comprising:
2 generating a slice correlation data structure for the packet based upon the storage location
3 of the first cell of each slice of the packet, and the packet identifier in each cell's descriptive
4 information;
5 responsive to the size of data having been changed as indicated in the indicator in the
6 system header, determining packet size adjustment; and
7 entering the slice correlation data structure as a single queue entry into a queue.
- 1 15. The method of claim 13 further comprising:

upon initiation of retrieval of the packet, generating a new packet identifier for the packet;
sending a dequeue message for each slice of the packet;
correlating each cell of the packet into packet form based on cell descriptive information including the packet position indicator and the slice position indicator; and
ordering the packet for transmission to an attached network based on the new packet identifier.

16. A multi-slice network processor system comprising:

a plurality of parallel processing slices, each processing slice comprising a lookup processing module and access to a storage sub-system;

a network data distribution and aggregation module for segmenting a packet received from a network into one or more packet slices, the network data distribution and aggregation module having a communication interface to each of the processing slices for communicating each packet slice;

each of the plurality of slices having a channel communication interface with the network interface over which each packet slice is directed to a destination processing slice across the network interface; and

a queuing module having an enqueueing communication interface and a dequeueing communication interface with each of the processing slices, the queuing module controlling the enqueueing and dequeueing of each of the packet slices, and determining the destination processing slice based on load balancing criteria.

17. The system of claim 16 wherein load balancing criteria includes no load balancing.

- 1 18. The system of claim 16 wherein the network interface is a switch fabric, and wherein each
2 channel communication interface comprises a port connection with the switch fabric.
- 1 19. The system of claim 16 wherein the storage sub-system includes a memory, the memory
2 storing at least one group of cells of a packet in a buffer; and a buffer manager, the buffer
3 manager maintaining a buffer correlation data structure for correlating one or more buffers of the
4 same packet slice, the buffer correlation data structure being stored in the memory.
- 1 20. The system of claim 16 wherein the queuing module includes a queuing memory space, the
2 queuing module maintaining a slice correlation data structure for correlating one or more slices
3 of the same packet slice in a single queue entry, the slice correlation data structure being stored
4 in the queuing memory space.
- 1 21. The system of claim 19 wherein the buffer manager comprises an ingress buffer manager
2 including an ingress buffer memory space for each processing slice, the ingress buffer memory
3 space for storing cells received from the respective processing slice, and an egress buffer
4 memory space for each processing slice, the egress buffer memory space for storing cells
5 received from the switch fabric for each respective processing slice.